

ADAPTIVE NOISE FILTERING AND EQUALIZATION FOR OPTIMAL HIGH SPEED MULTILEVEL SIGNAL DECODING

CROSS REFERENCE TO RELATED APPLICATIONS

The present application claims priority under 35 U.S.C. § 119(e) to U.S. Provisional Application Serial No. 60/396,065 entitled, "ADAPTIVE NOISE FILTERING AND EQUALIZATION FOR OPTIMAL HIGH SPEED SIGNAL DECODING," filed on July 15, 2002 in the name of Andrew Joo Kim et al. The entire contents of which are hereby incorporated by reference. This application is also related to U.S. Non-provisional Application Serial No. 10/108,598 entitled, "METHOD AND SYSTEM FOR DECODING MULTILEVEL SIGNALS," filed on March 28, 2002 in the name of Vincent Hietala et al., the entire contents of which are hereby incorporated by reference.

FIELD OF THE INVENTION

The present invention relates to communication systems and improving the received signal quality in a high-speed communications environment through the use of equalization. The improvement in signal quality affords gains in system performance such as increased data throughput capacity or reduced error rate. Specifically, the present invention relates to a method and system for improving the quality of a received signal by counteracting distortions introduced in signal generation, transmission, and reception.

BACKGROUND OF THE INVENTION

Network bandwidth consumption is rising at a rapid rate. Existing network capacity is marginally adequate and is expected, as of this writing, to soon be inadequate. Thus, there is a need to increase network bandwidth capacity. This increase in bandwidth can be achieved by increasing the symbol transmission rate to yield a corresponding increase in the data rate or by using advanced modulation techniques with higher spectral efficiency (i.e. techniques that communicate more than one information bit per symbol).

Regardless of the technique employed to achieve higher data throughput, the higher data throughput can place more stringent requirements on the fidelity of the signal communicated. Fidelity of the signal communicated can be hampered by signal degradation. Signal degradation can occur during signal generation and signal transmission. Signal degradation incurred in generating and transmitting a signal over a channel can largely be categorized as arising from two sources: (i) filtering of the signal and (ii) corruption from noise.

In classical communications (e.g. wireless or wireline communications), the noise component is commonly addressed by using optimal detection (i.e. matched-filtering followed by optimal thresholding). However, such a conventional approach often neglects the inter-symbol interference (ISI) associated with the filtering that occurs in the channel, i.e. that approach assumes that the noise is the dominant source of distortion. If the ISI is the dominant source of signal degradation, then the conventional approach is to equalize the channel, e.g. filter the received signal with an inverse filter prior to detection. The use of any one of these approaches in isolation may not improve signal fidelity since matched-filtering and equalization are often contradicting goals.

For example, equalization generally corresponds to high-pass filtering which, while removing ISI, increases the presence of high-frequency noise. A low-pass filter (LPF) is usually employed to the equalized signal in order to reduce the effect of the high-frequency noise but which also re-introduces ISI. Matched-filtering, on the other hand, is often low-pass in nature and thus frequently exacerbates the ISI in the signal in the process of reducing noise.

The separate application of matched-filtering and equalization can be characterized as “ad-hoc” because it does not consider the problem of noise mitigation and equalization in a combined framework, and thus, neglects the impact each has on the other.

There exist techniques in the conventional art which address noise mitigation and equalization in a common framework. In particular, the well-known Least-Mean Squares (LMS) based approaches minimize a distortion measure that captures the impact of both noise and ISI. Furthermore, these methods are adaptive in the sense that the settings of the filter are automatically adjusted to the optimal value. This adaptive feature is often necessary as the exact characteristics of the channel distortion and noise spectral content vary from installation to installation and also with time and temperature in some instances.

Unfortunately, the use of these traditional adaptive LMS-based control methodologies for high-data rate systems can be impractical due to data acquisition and processing difficulties. In particular, it can be economically impractical (and often technically infeasible) to (i) produce the analog-to-digital converters (ADC’s) capable of digitizing the signal at the required speed and resolution and (ii) produce a processor capable of handling the digitized data at the high speeds.

Therefore, there is a need in the art for an adaptive filtering approach that combines channel equalization and noise filtering. Another need exists in the art for a method and system for high speed digital communications that combines channel equalization and noise filtering in a single framework and that can account for the effects that equalization can have on noise

filtering, and vice-versa. Additionally, there is a need for such a method and system which is economically and technically practical for high-speed data systems.

SUMMARY OF THE INVENTION

A Signal Conditioning Filter (SCF) and a Signal Integrity Unit (SIU) can control elements that filter a digital (i.e. binary or multilevel) signal. A multilevel signal uses discrete amplitudes, which can change from one time interval to another, to convey information in each time interval. The simplest example of multilevel signaling is binary signaling where two amplitudes are used to represent a logical 0 or 1 (i.e. one bit of information). By using more levels, more information can be conveyed with each symbol or in each time interval. In some of the prior art and conventional art, the term “multilevel” conveys the use of more than two amplitude levels. To avoid any ambiguity, the term “digital signaling” will be used in this writing to refer to signaling with discrete amplitudes and time intervals, i.e. digital signaling that can include both binary and multilevel signaling.

The SCF and SIU form part of a method and system for equalizing and filtering a digital signal. This method and system for equalizing and filtering may be used in a variety of high-speed communications systems. Applications can include, but are not limited to, (i) electrical systems such as backplane, telecom, and datacom systems and (ii) optical systems such as long-haul, metro, and short-reach applications.

Regardless of the application, the method and system can process a received digital signal in the electrical domain prior to decoding. Thus, in optical systems, the method and system can be used either after photodetection in the receiver or prior to modulation in the transmitter.

The present invention can address the coupled problem of equalization and noise filtering in order to improve signal fidelity for decoding. Specifically, a received digital signal can be filtered in a manner to optimize the signal fidelity even in the presence of both significant (large magnitudes of) ISI and noise. Furthermore, the method and system of the present invention can be adaptive in the sense that filter coefficients can be continuously updated to reflect any time-varying changes in the system characteristics.

The present invention can provide an adaptive method that continuously monitors a signal fidelity measure. For example, monitoring the fidelity of a digital signal can be performed by external means such as a Signal Integrity Unit (SIU). A received signal $y(t)$ can be “conditioned” by application of a filter with an electronically adjustable impulse response

$g(t)$. A resulting output $z(t)$ can then be interrogated to characterize the quality of the conditioned signal.

This fidelity measure $q(t)$ can then be fed back to the SCF. Utilizing the signal fed back to the SCF, the response of the SCF can be adjusted to maximize the received fidelity measure. For the SIU, the signal fidelity measure can be directly associated with a decision error probability in a subsequent decoder with optimal decision thresholds. Combining the proposed approach with such a control system can balance (in a principled fashion) the trade-off between the degree to which ISI is corrected and noise is mitigated for optimal decoding.

The SCF can include a cascade of two or more tapped delay line filters with electronically controllable gain coefficients. The tap spacings of the two filters can be different in order to effectively combat both the effect of ISI which occurs on a large time scale and the effects of noise, jitter, and signal ringing which occur on a small time scale.

Using a cascade of two distinct filters can minimize the number of taps required to address both of these phenomena. The delay lines in these filters can include artificial transmission lines which can absorb the parasitic capacitance of the tap amplifiers. The tap amplifiers which vary the gain coefficients can be implemented using special Gilbert cell multipliers.

BRIEF DESCRIPTION OF THE DRAWINGS

Fig. 1A illustrates a digital signal receiver architecture according to one exemplary embodiment of the present invention.

Fig. 1B illustrates a signal integrity unit according to one exemplary embodiment of the present invention.

Fig. 2 illustrates an equalization architecture according to one exemplary embodiment of the present invention.

Fig. 3 illustrates a short-time filter and a long-time filter that form a signal conditioning filter according to one exemplary embodiment of the present invention.

Fig. 4 illustrates a short-time filter in the form of a tapped-delay line filter with small tap delays less than a symbol period of a digital signal according to one exemplary embodiment of the present invention.

Fig. 5. illustrates a long-time filter in the form of a tapped-delay line filter with delays substantially equal to the symbol period according to one exemplary embodiment of the present invention.

Fig. 6 illustrates a signal conditioning filter (SCF) forming part of a communications transmitter according to one exemplary embodiment of the present invention.

Fig. 7 is a functional block diagram illustrating a tapped-delay line filter according to one exemplary embodiment of the present invention.

Fig. 8 is a functional block diagram illustrating a tapped-delay line filter and one of the exemplary signal paths according to one exemplary embodiment of the present invention.

Fig. 9 is a functional block diagram illustrating a variable gain tap amplifier according to one exemplary embodiment of the present invention.

Fig. 10 is a circuit diagram illustrating a lumped LC delay line that has terminations provided on either end according to one exemplary embodiment of the present invention.

Fig. 11 is a circuit diagram illustrating a three-tap tapped-delay line filter having variable gain amplifiers described in Fig. 18 and delay blocks described in Fig. 19 according to one exemplary embodiment of the present invention.

Fig. 12 is a circuit diagram of a variable gain tap amplifier according to one exemplary embodiment of the present invention.

Fig. 13 is a circuit diagram of 5-tap tapped-delay line filter with artificial transmission lines according to one exemplary embodiment of the present invention.

Fig. 14 is a circuit diagram of an exemplary long-time delay artificial transmission line according to one exemplary embodiment of the present invention.

Fig. 15 is a circuit diagram of an exemplary three-tap tapped-delay line filter according to one exemplary embodiment of the present invention.

Fig. 16 is a graph that illustrates how the effective Q-factor varies with the initial iterations of the adaptive filter for an electrical multilevel signal according to one exemplary embodiment of the present invention.

Fig. 17 is a graph that illustrates how the effective Q-factor varies with the initial iterations of the adaptive filter for a multilevel signal going through 100 km of fiber according to one exemplary embodiment of the present invention.

Fig. 18 is a graph of an Eye-diagram for an unfiltered multilevel electrical drive signal that can be used as an input to the present invention.

Fig. 19 is a graph of an Eye-diagram for a filtered multilevel electrical drive signal according to one exemplary embodiment of the present invention.

Fig. 20 is a graph of an Eye-diagram for an unfiltered multilevel optical signal that can be used as an input to the present invention.

Fig. 21 is a graph of an Eye-diagram for a filtered multilevel optical signal according to one exemplary embodiment of the present invention.

Fig. 22 is a graph illustrating how the effective Q-factor varies with the initial iterations of the adaptive filter for a multilevel signal with inter-symbol interference according to one exemplary embodiment of the present invention.

Fig. 23 is a graph of an Eye-diagram for an unfiltered multilevel optical signal that can be used as an input to the present invention.

Fig. 24 is a graph of an Eye-diagram for a filtered multilevel optical signal according to one exemplary embodiment of the present invention.

Fig. 25 illustrates a signal conditioning filter (SCF) fabricated in GaAs as an integrated circuit that contains a variable gain input amplifier and a transit detector circuit for clock recovery according to one exemplary embodiment of the present invention.

Fig. 26 is a graph that illustrates a 5 Gbps unfiltered binary signal after transmission over 34" of copper trace in a backplane that can be used as input to the present invention.

Fig. 27 is a graph that illustrates the signal of Fig. 26 after equalization with the signal conditioning filter according to one exemplary embodiment of the present invention.

Fig. 28 is a graph that illustrates an unfiltered 10 Gbps (4-level 5G sym/s) signal after transmission over 150 m of multimode fiber that can be used as input to the present invention.

Fig. 29 is a graph that illustrates the multilevel signal of Fig. 28 after equalization with the signal conditioning filter according to one exemplary embodiment of the present invention.

DETAILED DESCRIPTION OF THE EXEMPLARY EMBODIMENTS

The present invention can address the problems of equalization and noise filtering in order to improve signal fidelity for decoding. Specifically, a received digital signal can be filtered in a manner to optimize the signal fidelity even in the presence of both large magnitudes of ISI and noise. Furthermore, the method and system of the present invention can be adaptive in the sense that filter coefficients can be continuously updated to reflect any time-varying changes in the channel behavior.

Filter Structure

Referring now to Fig. 1A, an equalization and filtering system 102 can comprise a variable gain amplifier 105, a signal conditioning filter (SCF) 100, a decoder or clock and data recovery (CDR) unit 120, and a signal integrity unit (SIU) 125. The CDR 120, SIU 125, and VGA 105 can form parts of a signal detection and fidelity characterization circuit 130 as will be discussed below with respect to Fig. 2. The exemplary signal conditioning filter (SCF) 100 can comprise two filters, a short-time filter 110 and a long-time filter 115 that are cascaded together.

According to one exemplary embodiment, each filter 110, 115 of the signal conditioning filter 100 can comprise a tapped-delay line or finite impulse response (FIR) filter with electrically controllable gain coefficients. Filtering is separated into these two stages, where the short-time filter 110 can comprise the first stage and the long-time filter can comprise the second stage. Each stage is designed to address a particular type of distortion using a relatively small number of taps.

As its name implies, the short-time filter 110 can be designed to mitigate degradations such as ringing, noise, and timing jitter that can occur on a relatively small time scale. Meanwhile, the long-time filter 115 is designed to remove signal artifacts (such as ISI) that can occur on a larger time scale.

Referring now to Fig. 1B, the signal integrity unit 125 can comprise a low pass filter (LPF) 132 and an analog-to-digital converter 135, a controller or processor 140, and a plurality of digital-to-analog converters (DACs) 145. After low-pass filtering by the LPF 132, a DC component of the sampled signal remains and is termed the event monitor voltage and is an analog probability estimate for the controlled reference voltage exceeding the received signal $y(t)$ where controlled reference voltage is generated by the digital-to-analog converters 145.

As mentioned above, the output of the LPF 132 can be fed to the analog-to-digital converter (ADC) 135. The ADC 135 may be characterized as a low-speed high-resolution ADC that measures the averaged event-detector output representing the cumulative distribution function (CDF) value. Specifically, the reference voltage is swept over a range of voltage levels while the ADC 135 samples the voltage from the filter 132 to produce an estimate of the CDF. The SIU 125 can set the reference voltage via the DAC 145 to a fixed value and then measures the averaged event detection output. The SIU 125 can then set the reference voltage to a different fixed value and measures another point of the CDF. This process is completed until the CDF curve is formed.

The ADC 135 feeds its output to a microcontroller 140. Microcontroller 140 can process the cumulative distribution function (CDF) to determine threshold voltage values for

signal decoding, receiver gain for the variable gain amplifier 105, and filter coefficients for the SCF 100. The microcontroller 140 is responsible for feeding the filter coefficients to the SCF for equalizing and filtering the received signal. Further details of the signal integrity unit 125 are discussed in the commonly owned U.S. Non-provisional Application Serial No. 10/108,598 entitled, "METHOD AND SYSTEM FOR DECODING MULTILEVEL SIGNALS," filed on March 28, 2002 in the name of Vincent Hietala et al., the entire contents of which are hereby incorporated by reference.

Referring now to Fig. 2, a first input to the signal conditioning filter (SCF) 100 can comprise a signal $y(t)$ that can include an unfiltered received signal while the output to the SCF 100 can comprise a filtered signal $z(t)$. The filtered signal $z(t)$ can be propagated into the signal detection and fidelity characterization circuit 130. As mentioned above, the signal detection and fidelity characterization circuit 130 can comprise a CDR 120 and SIU 125.

A first output of the signal detection and fidelity characterization circuit 130 can comprise the decoded data from the filtered signal $z(t)$. A second output of the signal detection and fidelity characterization circuit 130 can comprise a control signal $q(t)$ that forms a second input to the SCF 100.

Referring now to Fig. 3, as noted above with respect to Fig. 1, the SCF 100 can comprise a short-time filter 110 that receives the signal $y(t)$ discussed above with respect to Fig. 2. The short-time filter 110 manipulates the signal $y(t)$ according to a first function $g_s(t)$. The output from the short-time filter 110 can comprise a first filtered signal $y'-prime(t)$ that is fed as the input to the long-time filter 115.

The long-time filter 115 manipulates the signal $y'-prime(t)$ according to a second function $g_l(t)$. The output from the long-time filter 115 can comprise a second filtered signal or $z(t)$ as mentioned above with respect to Fig. 2 that is fed as the input to the CDR 120 or SIU 125 (not shown in Fig. 2).

Referring now to Fig. 4, the exemplary short-time filter 110 can comprise a tapped-delay line filter with amplifiers 405 and N delay elements 410 (each with delay δ) and $N+1$ amplifiers with gain coefficients α_n for $n=0,..,N$. The delay δ can be chosen to be small (relative to the symbol period T_0 of the signal) to permit the short-time filter 110 to perform at least one of three functions: (1) compensate for signal distortions (such as ringing) that can occur within a single symbol period; (2) effectively integrate over less than a symbol period to integrate out and reduce the noise; and (3) adjust the amount of signal smoothing to be large enough to reduce the effects of signal distortion and noise but to be short enough to be robust relative to the timing jitter in the system.

The short-time filter 110 can support a frequency resolution of $1/N\delta$. It is assumed that the delay δ is sufficiently small so that there is no aliasing in the short-time filter 110. Specifically, the short-time filter 110 has a frequency response that is periodic in frequency with period $1/\delta$. Any signal (or noise) energy at frequencies higher than $f=1/(2\delta)$ will distort the filtered signal as they overlap into adjacent spectral periods. Because of this distortion, it is recommended that the signal be pre-filtered with a passive analog low-pass filter (not shown) prior to the short-time filter 110. This pre-filtering may not be needed if any of the receiver components already sufficiently bandlimits the signal, as is often the case with circuit hardware and the high speeds considered.

The output of the short-time filter 110 can be written as

$$y'(t) = \alpha_0 y(t) + \alpha_1 y(t-\delta) + \dots + \alpha_N y(t-N\delta) \quad (1)$$

or equivalently as

$$y'(t) = [\alpha_0 + \alpha_1 + \dots + \alpha_N] y(t) - \alpha_1 [y(t) - y(t-\delta)] - \dots - \alpha_N [y(t) - y(t-N\delta)] \quad (2)$$

where the latter form explicitly conveys how the short-time filter 110 operates on the difference between the current sample and a sample from the past, i.e. each term can provide a first-order correction for the signal fluctuation within a symbol period. Furthermore, the coefficient on $y(t)$ in Eq. (2) provides the DC gain of the signal, i.e. the gain on the signal when the signal is already flat within a symbol period, and hence, all the differential terms are zero. As a reference, unity DC gain can be chosen by setting α_0 to

$$\alpha_0 = 1 - \alpha_1 - \dots - \alpha_N \quad (3)$$

or by alternatively normalizing the filter coefficients to sum to one.

Referring now to Fig. 5, the exemplary long-time filter 115 can comprise a tapped-delay line filter with M delay elements 410' (each with a delay equal to the symbol period T_0) and $M+1$ amplifiers 405 with gain coefficients a_m for $m=0,..,M$. The purpose of this long-time filter 115 is to remove ISI which occurs between symbols. Because the short-time filter 110 is capable of smoothing the signal over a significant portion of the symbol period, the long-time filter 115 need not worry about aliasing associated with the signal frequencies higher than the sampling rate $1/T_0$. If the short-time filter 115 were not present, then an anti-aliasing filter with a low-cutoff frequency may be needed.

Similar to Eqs. (1) and (2), the output of the long-time filter 115 can be written as

$$z(t) = a_0 y'(t) + a_1 y'(t-T_0) + \dots + a_M y'(t-MT_0) \quad (4)$$

or equivalently,

$$z(t) = [a_0 + a_1 + \dots + a_M] y'(t) - a_1 [y'(t) - y'(t-\Delta)] - \dots - a_M [y'(t) - y'(t-M\Delta)] \quad (5)$$

where again, a unity DC gain can be arbitrarily chosen by setting

$$a_0 = 1 - a_1 - \dots - a_N \quad (6)$$

or by scaling the coefficients to sum to one.

Those skilled in the art can extend the exemplary filters in Figures 4 and 5 to account for subsequent as well as preceding signal samples. For example, additional L taps could be added to the short-time filter 110 of Fig. 4 to change Eq. (2) to yield

$$y'(t) = [a_0 + \dots + a_{N+L}] y(t) - a_1 [y(t) - y(t-\delta)] - \dots - a_{N+L} [y(t) - y(t-(N+L)\delta)].$$

The signal sample $y(t-L\delta)$ can be interpreted as the sample to be decoded, i.e. where the decoding has been delayed by process by time $L\delta$. Because, the signal points $y(t)$ through $y(t-(L-1)\delta)$ precede $y(t-L\delta)$, the ISI associated from these symbols can be removed. Thus, the short-time filter 110 can now remove ISI originating from both preceding and succeeding signal samples. The long-time filter 115 may be similarly modified. This extension can be applied to both the short and long-time filters 110, 115 in practice as it can help improve equalization.

Filtering in the Transmitter

Referring now to Fig. 6, although described for use in a receiver, the SCF 110 may also be used in a transmitter 600 to maximize the fidelity of the transmit signal before channel noise is introduced which limits signal restoration in a corresponding receiver (not shown). Specifically, SCF 100 and SIU 125 modules may be placed in the transmitter 600 immediately before transmission circuitry as illustrated in Figure 6 where these modules can be used to minimize the effects of ringing, ISI, circuit impedance mismatches, and noise introduced by the transmitter.

Removing these distortions prior to transmission is advantageous because the noise introduced by the channel limits the degree to which the transmitter 600 distortions can be compensated by the receiver's SCF 100. Additionally, if a supervisory communications link is available between the transmitter 600 and receiver (not shown), then the fidelity signal $q(t)$ may be fed from the receiver (not shown) to the transmitter 600 over this link so that the SCF 100 in the transmitter not only compensates for transmitter distortions, but also pre-compensates for link distortions. This pre-compensation would be advantageous because it prevents the noise introduced by the link from limiting the degree to which the channel distortions can be compensated.

Filter Realization

The exemplary short- and long-time filters 110, 115 forming the SCF 100 can be implemented using an integrated tapped-delay line filter structure 100A with the basic structure illustrated in Figure 7. This exemplary filter structure 100A is functionally identical to the conceptual tapped-delay line filter diagrams of Figures 3 and 4 if τ is associated with δ for Figure 4 and T_0 for Figure 5.

The filter structure 100A shown in Figure 7 offers efficient “re-use” of the time delays for the various signal paths. This can be important since time delays are generally physically quite large and thus difficult to integrate. Therefore, it is desirable to make optimal use of all delay elements.

To understand this efficient “re-use” of the delays, shown in Figure 8 is the signal path for the third filter tap with gain coefficient g_3 . For this tap, by inspection, the total signal delay will be six times (X) $\tau/2$ or 3τ . Similarly, by observation, the delay of the signal for each tap g_i is i (times) X τ . Thus, the output signal $y(t)$ can be represented as follows:

$$y(t) = \sum_{k=0}^4 g_k x(t - k \tau), \quad (7)$$

which is identical to the form of Equations (1) and (4) discussed above.

A block diagram of an exemplary variable gain tap amplifier 900 for a tap filter is illustrated in Figure 9. In this exemplary embodiment, the gain constant GC is multiplied with the signal V_{in} to allow for both positive and negative gain coefficients. The input signal V_{in} is amplified/buffered, multiplied by the gain coefficient GC, and then output as a current source I_{out} . The input amplifier/buffer 902 is designed to have high input impedance as to not disturb the input delay line. The output of the amplifier 900 is also designed to be high impedance as to similarly not disturb the output delay line. In practice, a significant parasitic capacitance remains at both the input and output of the amplifier 900, but as will be seen below, these parasitics can be absorbed into the delay lines.

For one exemplary embodiment of the present invention, a delay element can comprise a simple *LC* (inductor-capacitor) delay line. However, those skilled in art will recognize that other delay elements can be used and are not beyond the scope and spirit of the present invention.

A representative *LC* delay line 1800 is illustrated in Fig. 10. This delay line 1000 forms a high order low-pass filter function with a relatively constant delay over the passband. Those skilled-in-the-art will recognize that a variety of filter design methodologies could be employed to construct higher performance delay structures, but the one shown in Fig. 10 was used for

simplicity. It should also be realized that the lumped elements could be realized by distributed elements or in fact a variety of different delay elements could be similarly used.

For the simple delay structure shown, assuming constant L and C_l , the delay, τ is approximately:

$$\tau \approx N_{LC} \sqrt{LC_l} \quad (8)$$

in which N_{LC} is the number of LC pairs and the filter's termination resistance, R_o is:

$$R_o \approx \sqrt{\frac{L}{C_l}} \quad (9)$$

The end of each delay element needs to be terminated with R_o in parallel with a capacitor of approximately $C_l/2$ for proper operation.

Referring now to Fig. 11, an exemplary embodiment of a three tap tapped-delay line filter 100C using the variable gain tap amplifier of Fig. 9 and the lumped LC delay element shown in Fig 10 is illustrated. Several important design aspects of this filter structure 100C can now be explained.

First, it is important to realize that the cascade of delay elements 1102 and circuitry on the left-hand side of the circuit 100C is designed to form a well controlled delay line. This will be referred to as the "input delay line" 1102. Similarly, on the right-hand side of the circuit 100C comprising the two delay elements and associated circuitry in this simple exemplary embodiment, will be referred to as the "output delay line" 1104.

Both the input and output delay lines 1102, 1104 are designed to have minimal signal attenuation and reflections in order to maintain good signal fidelity. As such, as already mentioned, the inputs of the amplifiers g were designed with high impedance and the remaining input capacitance is "absorbed" into the loading capacitance of the input delay line 1102.

This absorbed capacitance can be seen in Fig. 11 by the reduced values of loading capacitance ($C_l - C_{in}$). Additionally, the input and the output of the input delay lines 1102, 1104 are carefully terminated in the characteristic resistance (R_o) of the delay elements. Similarly, the output delay line 1104 absorbs the output parasitic capacitance of the amplifiers and is terminated with the appropriate resistance (R_o). The output of each amplifier g launches a signal in both directions in the output delay line 1104 and thus it can be appreciated why proper termination can be critical for proper filter operation.

Figures 12 through 15 illustrate an exemplary circuit embodiment of two tap tapped-delay line filters. The designs are fully differential, but one skilled-in-the-art will realize that the concepts discussed above are equally applicable to differential or single-ended designs. The

particular exemplary embodiments illustrated in these figures are for integration in a GaAs HBT process. One skilled-in-the-art will realize that other device technologies could similarly be applied.

Specifically, Fig. 12 illustrates the circuit diagram of an exemplary variable gain tap amplifier 1200. The input is first buffered by high input-impedance emitter follower amplifiers X16 and X17. The output of the emitter follower amplifiers drives a lower pair of a standard Gilbert cell multiplier circuit (also called a Gilbert cell mixer and an XOR gate) comprising X14, X13, X9, X10, X11 and X12. The bases of the upper cross connected differential pairs are held at the desired DC gain constant. Since this circuit is an effect 4-quadrant multiplier, the gain coefficient can be negative or positive.

The R8, R10 and R9, R11 resistor divider networks scale the input drive voltage so that desired coefficient control range is achieved (approx. ± 1 V in this case). The output of the Gilbert cell multiplier circuit is terminated by the output delay line. The Gilbert cell multiplier will usually see a real load resistance of $R_o/2$. The output will have a significant output capacitance due the collector capacitances of X9, X10, X11, and X12, but as discussed above, this capacitance can be absorbed into the output delay line.

Referring now to Fig. 13, the design for an exemplary 5-tap $\tau=T_0/5$ short-time tapped-delay line filter 100D that is based on artificial transmission lines is illustrated. For this filter 100D, it was found that only one LC pair was required for each tap (2nd order low-pass filter). Each delay element can provide 37 (pico-seconds) ps of delay (1/2 of 1/5 period at 2.7 Gbps). The inductors used were integrated spiral inductors with a nominal inductance of 2.7 nH. The actual loading capacitor values were initially selected as indicated above, but optimized for the best frequency response in the actual design. A series resistance R1304 was added to the loading capacitors C1304 to lower the phase peaking at the stopband edge. An additional shunt resistance R1306 was added to help counteract the effect of the inductor loss. Inductor loss makes the delay lines characteristic resistance complex. The addition of an appropriate shunt resistance can cancel this effect and make the characteristic impedance real at a target frequency. The input and output of the filter 100D are buffered by amplifiers 1302 to isolate the circuit from external influences.

For the long-time filter 115 (3-tap $\tau=T_0$), a six stage LC filter was required to obtain the necessary delay of 185 ps. Referring now to Fig. 14, this figure illustrates an exemplary circuit of an individual long delay element 1400. Four of these long delay elements 1400 were used in the exemplary filter embodiment 100E illustrated in Fig. 15. The exemplary filter embodiment 100E of Fig. 15 can be referred to as a 3-tap long-time tapped-delay line filter.

Coefficient Adaptation Algorithm

The gains of the filter tap amplifiers g are adjusted to maximize a signal fidelity measure $q(t)$ (provided by the SIU 125 for example). The general idea is to slightly perturb the coefficients and observe the effect on $q(t)$. Doing this allows a set of coefficient values to be determined that locally maximizes the signal fidelity. The approach continually perturbs the coefficients to survey the signal fidelity's dependence on the coefficients and tune the SCF 100 to the optimal values in an adaptive fashion. The method is presented for the case of the short-time filter, but the approach is equally applicable to the long-time filter. Alternatively, the cascade of the two filters can be viewed as a single filter whose coefficients are adapted by the following method as will be discussed below.

The intuition in the preceding paragraph can be made precise by restating it as the following optimization problem

$$\alpha^* = \arg \max_{\alpha} \{q(t)\} \quad (10)$$

where α denotes the vector of adjustable filter coefficients and $q(t)$ quantifies the fidelity of the signal (such as a measure of the signal's Q-factor as produced by the SIU[4]). Eq. (10) is solved via an empirical minimization of $q(t)$. Towards this end, a coordinate descent algorithm can be used to find the local maximum by searching over each coordinate while the others are held fixed (i.e. perturbing one of the filter coefficients $\alpha_0, \alpha_1, \dots, \alpha_N$ at a time). One skilled in the art will realize that other numerical optimization techniques (such as gradient, Newton, and conjugate methods) may also be used to solve Eq. (10).

In some contexts, simply maximizing the signal fidelity may not be sufficient to provide good SCF performance. In particular, there is the possibility of a null-space of solutions. For example, consider the task of short-time equalization when there is no ringing and noise on the received signal to compensate. An intuitive solution would be to have $\alpha_0=1$ and $\alpha_n=0$ for all other n , i.e. perform no filtering. However, an equally valid solution would be, $\alpha_0=1$, $\alpha_1=A$, $\alpha_2=-A$, and $\alpha_n=0$ for all other n where A can be any value (including arbitrarily large values). Other convoluted (but still valid) sets of coefficients can also be obtained. The drawbacks of such non-intuitive solutions are

1. They obscure the identifiability of the coefficients which are actually helping to improve signal fidelity. For instance, in the above example, α_1 and α_2 effectively cancel each other and do nothing to reduce ringing or noise.

2. They reduce the robustness of the system to changes in the channel characteristics. In particular, if the channel characteristics change such that the null-space "moves" (e.g. a little ringing is introduced), then α_1 and α_2 will no longer cancel each other and a very erratic signal will result until the coefficients are relearned to remove such artifacts.

Motivated by such problems, there may be the desire to guide the values of the filter coefficients α_n , in addition to maximizing the signal fidelity. To achieve such a result, a regularization penalty can be imposed in the objective function. For example, solving the following optimization problem could be chosen in place of Eq. (10):

$$\alpha^* = \arg \max_{\alpha} \left\{ q(t)^2 - \gamma \|\alpha - \beta\|_2^2 \right\} \quad (11)$$

where γ and β are regularization parameters. The parameter β is a nominal value for α . Specifically, β is the value one would expect α^* to be in the absence of any channel distortion. While β represents the nominal value of α^* , γ represents confidence in β and determines how strongly α^* is driven towards β . Note that if no biasing is desired, then γ should be set to zero, and Eq. (11) reduces to Eq. (10). Thus, Eq. (10) can be seen as a specific case of Eq. (11). One skilled in the art will realize that a variety of penalty functions can be used in place of the 2-norm in Eq. (11).

As previously stated, the long-time filter coefficients can similarly be adapted according to Eq. (11), i.e. it can be chosen that

$$\alpha^* = \arg \max_{\alpha} \left\{ q(t)^2 - \gamma \|\alpha - b\|_2^2 \right\} \quad (12)$$

where b is the nominal value for the long-time filter coefficients. The nominal situation again corresponds to when the channel introduces no signal distortion. For example, if one generally believes the channel introduces no ISI, then b should generally resemble a Euclidean basis vector, i.e. a vector with a single 1 and the rest zeros. That is, the nominal long-time filter should not do much filtering.

Automatic Gain Control

Thus far, the special case of no DC amplification on the filtered multilevel signal has been considered. A non-unity gain could be selected. And if a variable gain amplifier (VGA) module 105 is to lead or follow the SCF 100, then the control of the VGA module 105 may be incorporated into the SCF's 100 control. In particular, if the desired gain A_{VGA} can be obtained as an external signal (perhaps determined in the SIU 125, then the SCF 100 can control the gain by simply scaling all of the coefficients α_n (or a_n) by A_{VGA} . This is simply a mathematical

operation that can be performed after the calculation of filter coefficients in Eq. (11) or (12). Alternately, the gain control of the VGA module 105 can be adjusted as needed by the SIU.

Simulation Results

To demonstrate results of the proposed filtering method, the inventors completed a MATLAB simulation as applied to real data of a 4-PAM signal in an optical communications system.

First, a system dominated by short-time effects, i.e. ringing and noise, is examined. Specifically, an electrical drive signal and a received signal after 100km of fiber is considered. The progression of the fidelity measure as the filter adaptation algorithm iterates is shown in Figures 16 (for the electrical drive signal) and 17 (for the received optical signal). Specifically, the graph 1600 in Figure 16 conveys the value of the signal fidelity measure as the filter coefficients are being adjusted. From the plot 1602, it can be seen that the fidelity of the signal is improved as the coefficients are adjusted. The numerical improvement conveyed by plot 1602 can be further appreciated by considering the eye-diagrams in Figures 18 and 19. Figure 18 shows the eye diagram for the signal before any filtering, i.e. at iteration 0 in graph 1600. Figure 19 shows the eye-diagram of the signal after filtering on iteration 7 in graph 1600. Clearly, the amount of ringing on the signal has been reduced through the use of filtering providing a better quality signal.

Figure 17 is analogous to Figure 16, except that Figure 17 is for a photodetected optical signal after 100km of optical fiber. Specifically, graph 1700 in Figure 17 shows the progression of the fidelity measure as the filter coefficients are adjusted in each iteration for this optical system. Figures 20 and 21 are analogous to Figures 18 and 19. In particular, Figure 20 shows an eye-diagram of the photodetected optical signal before any filtering. Clearly, the signal is of poor quality as exhibited by the nearly closed eye-openings 2002C. Figure 21, however, shows very wide eye-openings 2002D that are produced after filtering the signal with the described invention.

In addition to the short-time distortion noise dominated data sets in Figures 16 through 22, another data set in which the signal is predominantly degraded by ISI is examined. Again, the context is an optical system with 100km of fiber, but contrary to the system used for Figures 8, 11, and 12, a different photodetector is used which is less vulnerable to noise, thus leaving ISI as the dominant source of signal degradation. The evolution of the signal fidelity is shown in Figure 22 from which we see a significant improvement in signal quality (almost doubling the fidelity measure). The eye-diagrams associated with the SCF are shown in Figures 23 (without

filtering) and 24 (with filtering). The increase in the eye-opening from before filtering 2302A to after filtering 2302B clearly conveys the improvement in signal quality afforded by the SCF.

Laboratory Results

An exemplary IC implements the filter embodiment discussed above, as illustrated by the circuit 2500 of Fig. 25. The circuit 2500 was fabricated in a GaAs HBT process. The short-time filter 2505 comprises five coefficient amplifiers 2510 and eight delay elements 2515 to realize the 5-tap tapped-delay line filter illustrated in Figure 22. As previously described, a pair of $\pi/2$ delay elements 2515 are used to produce the desired delay of $\tau(\pi/2)$ from the input side and $\pi/2$ from the output side).

Meanwhile, the long-time filter 2520 comprises three coefficient amplifiers 2510 and four delay elements 2525 to realize the 3-tap long-time filter illustrated in Figure 24. Coefficient amplifiers 2510 are Gilbert cell multipliers with identical topologies, although specific circuit parameters may vary as each multiplier is optimized according to the surrounding circuit.

The improved output provided by the SCF circuit 2500 of Figure 25 is illustrated by Figures 26 and 27. Figure 26 illustrates a received signal from a backplane communications system transmitting a binary signal at a data rate of 5Gbps over a 34" copper trace. Specifically, Fig. 26 illustrates the condition of the received electrical binary signal in the form of an eye-diagram 2600. Clearly, no reliable data can be recovered from this received signal as there is no eye-opening.

Meanwhile, referring now to Fig. 27, this figure shows an eye-diagram 2700 for the same received signal illustrated in Fig. 26 but being processed by the SCF 100 and its corresponding control method. The eye-openings 2705 are now clearly visible affording reliable communications which was unachievable without equalization.

Referring now to Figs. 28 and 29, these figures demonstrate the improvement provided by the SCF 100 in an optical communications system. The system channel now includes optical components. Specifically, the electrical data signal is converted to an optical signal with a VCSEL (vertical cavity surface emitting laser), transmitted over 150m of multimode fiber (MMF), and converted back to an electrical signal at the receiver with a photodetector.

Fig. 28 illustrates an eye diagram 2800 where the condition of the received electrical signal when a 4-level 5Gsym/s (yielding a 10 Gbps data rate) is transmitted. Again, there are no visible eye-openings in the received signal and thus data cannot be reliably recovered.

Meanwhile, in Fig. 29, an eye-diagram 2900 for the same received signal is illustrated but after filtering with the SCF 100 of one exemplary embodiment of the present invention. The four signal levels are now apparent from the eye-openings 2905 which means that data can be recovered from this received and filtered signal.

Therefore, the present invention provides an adaptive filtering approach that combines channel equalization and noise filtering. The method and system of the present invention can easily support high speed digital communications which combines channel equalization and noise filtering in a single framework. The method and system of the present invention can account for the effects that equalization can have on noise filtering, and vice-versa. Furthermore, the method and system are adaptive in nature and have a practical means of implementation for high-speed data communications systems.

It should be understood that the foregoing relates only to illustrate the embodiments of the present invention, and that numerous changes may be made therein without departing from the scope and spirit of the invention as defined by the following claims.